

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 July 2001 (12.07.2001)

PCT

(10) International Publication Number
WO 01/50708 A2

(51) International Patent Classification⁷: **H04L 29/06**

(74) Agents: **BIRD, William, E. et al.**; Bird Goën & Co, Vilvoordsebaan 92, B-3020 Winksele (BE).

(21) International Application Number: **PCT/EP00/13392**

(22) International Filing Date:
29 December 2000 (29.12.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
99204623.5 31 December 1999 (31.12.1999) EP
09/490,398 24 January 2000 (24.01.2000) US

(71) Applicant (for all designated States except US): **ACR INC.** [US/US]; 278-A New Dorp Lane, Staten Island, NY 110306 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.


(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— Without international search report and to be republished upon receipt of that report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: A SERVER MODULE AND A DISTRIBUTED SERVER-BASED INTERNET ACCESS SCHEME AND METHOD OF OPERATING THE SAME



(57) Abstract: A wide area data carrier network is described comprising: one or more access networks; a plurality of server units housed in a server module and installed on said wide area data carrier network so that each server module is accessible from the one or more access networks, the server module being adapted so that it may be located at any position in the wide area network; and an operations centre for management of the server module, the server module being connected to the operations centre for the exchange of management messages through a network connection. The server module comprises at least one server card insertable in the server module, the server card having a central processing unit and at least one rewritable, non-volatile disc memory device mounted on the card. Preferably a local management system is provided in each server module capable of receiving a command from the operations centre and for executing the command to modify a service performance of at least one server unit.

A SERVER MODULE AND A DISTRIBUTED SERVER-BASED INTERNET ACCESS SCHEME AND METHOD OF OPERATING THE SAME

The present invention relates to the provision of server capacity in a wide area digital telecommunications network, in particular on any system which uses a protocol such as the TCP/IP set of protocols used on the Internet. The present invention also relates to a multi-server device which may be used in the digital telecommunications network in accordance with the present invention. The present invention also relates to a computing card for providing digital processing intelligence.

TECHNICAL BACKGROUND

A conventional access scheme to a wide area digital telecommunications network 1 such as the Internet is shown schematically in Fig. 1, which represents an IT centric Application Service Provider (ASPR) architecture. All servers 18 are deployed in a central data centre 10 where a data centric infrastructure is created to install, host and operate the ASPR infrastructure. Conventional Telecom Operators and Internet Service Providers are becoming interested in becoming Application Service Providers, in order to have a new competitive advantage by providing added value services in addition to their existing bearer services provided to telephone subscribers.

Application provisioning through IP networks, such as the Internet, is an emerging market. Service Providers in general have to provision application services in their network infrastructure. For this purpose, IT data centres 10 are conventionally used. Typically, the application servers 18 on which the applications offered are stored are located at the data centre 10 as well as some centralised management functions 16 for these application servers 18. Access is gained to these servers 14 via a "point of presence" 12 and one or more concentrators 14. A customer 11 dials a telephone number for a POP 12 and is connected to an Internet provider's communications equipment. Using a browser such as Netscape's Navigator™ or Microsoft's Explorer™ a session is then typically set up with an application server 18 in the remote data centre 10. Typically, a protocol stack such as TCP/IP is used to provide the transport layer and an application program such as the abovementioned browser, runs on top of the transport layers. Details of such protocols are well known to the skilled person (see for example, "Internet: Standards and Protocols", Dilip C. Naik, 1998

**(19) World Intellectual Property Organization
International Bureau**



(43) International Publication Date
12 July 2001 (12.07.2001)

PCT

(10) International Publication Number
WO 01/50708 A2

(51) International Patent Classification⁷: H04L 29/06

(74) Agents: BIRD, William, E. et al.; Bird Goën & Co, Vilvoordsebaan 92, B-3020 Winksele (BE).

(21) International Application Number: PCT/EP00/13392

(22) International Filing Date: 29 December 2000 (29.12.2000)

(25) Filing Language: English

(26) **Publication Language:** English

(30) Priority Data:

99204623.5	31 December 1999 (31.12.1999)	EP
09/490,398	24 January 2000 (24.01.2000)	US

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(71) **Applicant** (for all designated States except US): ACR INC. [US/US]; 278-A New Dorp Lane, Staten Island, NY 110306 (US).

Published:


— Without international search report and to be republished upon receipt of that report.

(72) Inventors; and

(75) Inventors/Applicants (for US only): DUJARDIN, Serge [BE/BE]; St. Truidersteenweg 178, B-3500 Hasselt (BE). **PARI, Jean-Christophe** [FR/FR]; Villa Samaria, 571 chemin de la Lauve, F-83700 Saint-Raphael (FR).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: A SERVER MODULE AND A DISTRIBUTED SERVER-BASED INTERNET ACCESS SCHEME AND METHOD OF OPERATING THE SAME



(57) Abstract: A wide area data carrier network is described comprising: one or more access networks; a plurality of server units housed in a server module and installed on said wide area data carrier network so that each server module is accessible from the one or more access networks, the server module being adapted so that it may be located at any position in the wide area network; and an operations centre for management of the server module, the server module being connected to the operations centre for the exchange of management messages through a network connection. The server module comprises at least one server card insertable in the server module, the server card having a central processing unit and at least one rewritable, non-volatile disc memory device mounted on the card. Preferably a local management system is provided in each server module capable of receiving a command from the operations centre and for executing the command to modify a service performance of at least one server unit.

WO 01/50708 A2

A SERVER MODULE AND A DISTRIBUTED SERVER-BASED INTERNET ACCESS SCHEME AND METHOD OF OPERATING THE SAME

5 The present invention relates to the provision of server capacity in a wide area digital telecommunications network, in particular on any system which uses a protocol such as the TCP/IP set of protocols used on the Internet. The present invention also relates to a multi-server device which may be used in the digital telecommunications network in accordance with the present invention. The present invention also relates to a computing card for providing digital processing intelligence.

10

TECHNICAL BACKGROUND

A conventional access scheme to a wide area digital telecommunications network 1 such as the Internet is shown schematically in Fig. 1, which represents an IT centric Application Service Provider (ASPR) architecture. All servers 18 are deployed
15 in a central data centre 10 where a data centric infrastructure is created to install, host and operate the ASPR infrastructure. Conventional Telecom Operators and Internet Service Providers are becoming interested in becoming Application Service Providers, in order to have a new competitive advantage by providing added value services in addition to their existing bearer services provided to telephone subscribers.

20 Application provisioning through IP networks, such as the Internet, is an emerging market. Service Providers in general have to provision application services in their network infrastructure. For this purpose, IT data centres 10 are conventionally used. Typically, the application servers 18 on which the applications offered are stored are located at the data centre 10 as well as some centralised management functions 16
25 for these application servers 18. Access is gained to these servers 14 via a "point of presence" 12 and one or more concentrators 14. A customer 11 dials a telephone number for a POP 12 and is connected to an Internet provider's communications equipment. Using a browser such as Netscape's Navigator™ or Microsoft's Explorer™ a session is then typically set up with an application server 18 in the remote
30 data centre 10. Typically, a protocol stack such as TCP/IP is used to provide the transport layer and an application program such as the abovementioned browser, runs on top of the transport layers. Details of such protocols are well known to the skilled person (see for example, "Internet: Standards and Protocols", Dilip C. Naik, 1998

Microsoft Press).

These IT centric data centres 10 may be suitable within the confines of a single organisation, i.e. on an Intranet, but in a network centric and distributed environment of telecom operators and Internet Service Providers such a centralised scheme can
5 result in loss of precious time to market, in increased expense, in network overloads and in a lack of flexibility. From an infrastructure point of view IT data centres 10 are very different from Telecom Centres or POPs 12, 14. The executed business processes that exploit an IT data centre are very different from business processes that have been designed for operating telecom and Internet wide area environments. It is expensive to
10 create a carrier class availability (99.999%) in an IT centric environment. Maintaining an IT environment (Operating Systems and applications) is very different from maintaining a network infrastructure for providing bearer services because of the differences in architecture. IT centric environments do not scale easily. Where it is planned that hundreds of potential subscribers will access the applications a big
15 "mainframe" system may be installed. Upgrading from a small to a medium to a large system is possible but this is not graceful – it implies several physical migrations from one system to another. Telecom Networks support hundreds of thousands of customers and do this profitably. To support this kind of volume it is difficult to provide and upgrade IT centric architectures in an economic manner. Since all the application
20 servers 18 are centrally deployed, all of the subscribers 11 (application consumers) will connect to the centre of the network 1. Typically the HQ where most of his IT resources are based. By doing this, network traffic is forced from the network edges into the network centre where the application servers are installed. Then, all the traffic has to go back to the network edge to deliver the information to the networked
25 application client. The result is that expensive backbone bandwidth usage is not optimised and packets are sent from edge to centre and back only because the location of the application servers.

IT centric application providers generally have two options for setting up the provisioning platform in the data centre 10. Either a dedicated server platform (i.e. one
30 application per server) or a shared server (i.e. multiple applications per server) is set-up. As an example, one server could be provided for per e-merchant wishing to run an e-shop on the server or multiple e-merchant shops could be set up on a single server. Setting up, maintaining, expanding and adapting business or network applications that

integrate many players (suppliers, partners, customers, co-workers or even children wanting to play "Internet games") into a common web-enabled chain is becoming increasingly complex. Such networked applications often require sophisticated multi-tiered application architectures, a continuously changing infrastructure, 24 hour, seven days a week availability, and the ability to handle rapid and unpredictable growth. While individual large companies often have a highly skilled IT personnel department and financial resources to meet these demands, many Service Providers cannot provide such services. For many telecom operators or Internet service providers that are preparing to become an application service provider, the only viable option is to host applications in a specially created and centralised data centre 10 where additional specially trained staff can be employed economically. Only when this "infrastructure" is complete, can applications be delivered via the Internet to the "application consumers" 11.

A theoretical advantage of this conventional approach is that all resources are centralized so that resources can be shared and hence, economy of scale can be achieved for higher profits and a better quality of service. The advantage is theoretical because the ASPR is facing a potential "time bomb" in the cost of operations as their subscriber population explodes. Also, the initial price tag per user that comes along with shared (fault tolerant) application servers is very high in comparison to the infrastructure cost per user in telecom environments.

As can be seen from Fig. 1, and independent of the topology of the network, data from all subscribers 11 accessing their subscribed applications through the designated POP 12, will transit the network 1 until it has reached the application data centre 10. High capacity network pipes need to be provisioned everywhere in the network 1 in order to guarantee the throughput required to obtain acceptable application performance. Even bigger links need to be provisioned around the application data centre 10 itself to guarantee acceptable application performance. It is difficult for network planners to provision the network bandwidth without knowing exactly which future applications, requiring unknown bandwidth will be accessed from any or all of POP's by a yet undefined number of subscribers simultaneously.

Difficulties have already been reported. News items on television can cause a rush to access certain WEB sites. If thousands of people do this at the same time (e.g. as caused by a pop concert sent live over the Internet, or a special and very attractive

offer on the Internet of limited duration) the present infrastructure cannot deal with the data flow and many cannot access the site.

The problem can become a vicious circle - first subscriptions are sold for services, the applications are then provisioned to provide the services and as this type of business expands the network pipes have to be upgraded. This has a direct and negative consequence that the users in the start-up phase or at some later time will have unacceptable and/or unpredictable application response times and will find the application performance behaviour unsatisfactory. An alternative is to forecast the application provisioning success and to invest accordingly into the network and data centre infrastructure based on the commercial forecasts. This places the financial risk with the provider since there are so many "unknown" variables.

Another disadvantage of IT centric shared server architecture shown in Fig. 1 is security and the maintenance of a secure environment. One of the first rules in security is to keep things simple and confinable. The system is preferably limited to a confinable functionality that can be easily defined, maintained and monitored. Implementing shared network application servers that will provision hundreds of different applications for several hundred thousand of application users is, from a security policy point of view, not realistic without hiring additional security officers to implement and monitor the security policy that has been defined.

The IT centric way of implementing application provisioning may be satisfactory in the beginning but it does not scale very well either from a network/traffic point of view, or from an application maintenance point of view, or from a security point of view.

Another difficulty with modern application software is that few users can adequately use all the functionality provided. This is left to experts. This has resulted in large IT departments to maintain both the software and the hardware of work stations, personal computers and the Local Area networks to which they are attached. The size and cost of these departments adds considerable cost to any operation and is prohibitive for small and medium size enterprises. Loss of LAN connectivity can cripple the operation of a company if it lasts for a few hours during which time no fax can be sent, no document can be printed unless standalone devices are provided as a back-up. There is a requirement to allow economic provisioning and maintenance of word processing, scheduling and financial applications in Small- and Medium-sized

Enterprises (SME).

WO 98/58315 describes a system and method for server-side optimisation of data delivery on a distributed computer network. User addresses are assigned to specific delivery sites based on analysis of network performance. Generalised
5 performance data is collected and stored to facilitate the selection of additional delivery sites and to ensure the preservation of improved performance. However, there is no disclosure of how to manage a plurality of servers on an individual basis from a remote location.

US 5,812,771 describes a system for allocating the performance of applications
10 in an networking chassis among one or more modules in the chassis. This allows allocation of applications among the network modules within the chassis. However, the management system cannot carry out commands received from a remote operations centre to modify the service performance of an individual network module.

It is an object of the present invention to provide a communications network, a
15 method of operating the same and network elements which enable the service provider provision applications, such as e-commerce, web hosting, intranet mail, distant learning applications etc. in a fast, easy and cost effective way.

It is an object of the present invention to provide a communications network, a method of operating the same and network elements with which the business risks are
20 lower than with conventional systems.

It is an object of the present invention to provide a communications network, a method of operating the same and network elements which can be gracefully upgraded without either a high initial outlay or a lot of major network upgrades later.

It is an object of the present invention to provide a communications network, a
25 method of operating the same and network elements with improved flexibility and response times.

It is an object of the present invention to provide a server module which can be used as a network element and a method of operating the same which provides high security of data and application programs as well as a high security of availability.

30 It is an object of the present invention to provide a server module which can be used as a network element and a method of operating the same which is easy to maintain by non-engineer grade staff.

SUMMARY OF THE INVENTION

The present invention may provide a wide area data carrier network comprising: one or more access networks; a plurality of server units housed in a server module and installed in said wide area data carrier network so that each server module is accessible from the one or more access networks, and an operations centre for remote management of the server module, the server module being connected to the operations centre for the exchange of management messages through a network connection. Preferably each server module includes a management system local to the server module for managing the operation of each server unit in the module. The operations centre manages each server unit via the local management system. The local management system may be a distributed management system which is distributed over the server units in one module but more preferably is a separate management unit. The local management system is capable of receiving a command from the operations centre and executing this command so as to modify the service performance of at least one of the server units. Modifying the service performance means more than just reporting the state of a server unit or selecting a server unit from the plurality thereof. That is the local management system is capable of more than just monitoring the server units.

The local management system may also include a load balancing unit. This load balancing unit may be used for load balancing applications running on the server units, e.g. an application may be provided on more than one server unit and the load on each server unit within the group running the application is balanced by the load balancing unit; or a load balancing unit may be used for load balancing network traffic, e.g. to balance the loads on proxies used to transmit received messages to the relevant server unit. The server units may be active servers (rather than passive shared file message stores). The network connections to the server module may be provided by any suitable connection such as an interprocess communication scheme (IPC), e.g. named pipes, sockets, or remote procedure calls and via any suitable transport protocol, e.g. TCP/IP, etc. The management function may include at least any one of: remote monitoring of the status of any server unit in a module, trapping alarms, providing software updates, activating an unassigned server module, assigning a server module to a specific user, extracting usage data from a server module or server unit, intrusion detection (hacker detection). Preferably, each server unit is a single board

server, e.g. a pluggable server card. Preferably, each server unit includes a central processor unit and a secure memory device for storing the operating system and application programs for running the server unit. A rewritable, non-volatile storage device such as a hard disk is provided on the server unit. The server unit is preferably adapted so that the rewritable, non-volatile storage device contains only data required to execute the application programs and/or the operating system program stored in the secure memory device but does not contain program code. In particular the CPU is preferably not bootable via the rewritable, non-volatile storage device. Preferably, the server module is configured so that each server unit accesses the administration card at boot-up to retrieve configuration data for the respective server unit. In particular, the server unit retrieves its internal IP address used by the proxy server card to address the server unit. Preferably, each server unit is mounted on a pluggable card. The server card is preferably plugged into a backplane which provides connections to a power supply as well as a data connection to other parts of the server module connected in the form of a local area network.

The present invention also includes a method of operating a wide area data carrier network having one or more access networks comprising the steps of: providing a plurality of server units housed in a server module in said wide area data carrier network so that each server module is accessible from the one or more access networks; and managing each server unit of the server module remotely through a network connection to the server module via the local management system. Preferably each server unit of a server module is managed by a management system local to the server module. The remote management of each server unit is then carried out via the local management system. Local management includes the steps of receiving a command from the operations centre and executing this command so as to modify the service performance of at least one of the server units. Modifying the service performance means more than just reporting the state of a server unit or selecting a server unit from the plurality thereof. That is the local management includes more than monitoring the server units.

The local management may also include a load balancing step. This load balancing step may balance the load of applications running on the server units, e.g. an application may be provided on more than one server unit and the load on each server unit within the group running the application is balanced; or a load balancing step may

balance the load of network traffic, e.g. to balance the loads on proxies used to transmit received messages to the relevant server unit.

The present invention also includes a server module comprising:
a plurality of server cards insertable in the server module, each server card providing
5 an active server, e.g. a network server. Each server card is preferably a motherboard
with at least one rewritable, non-volatile disc memory device mounted on the
motherboard. The motherboard includes a central processing unit and a BIOS memory.
An Input/Output (I/O) device is preferably provided on the card for communication
with the central processing unit, for example a serial or parallel port. At least one local
10 area network interface, is preferably mounted on the server card, e.g. an Ethernet™
chip. Preferably, the operating system for the central processing unit and optionally at
least one application program is pre-installed in a solid state memory device.
Preferably, the program code for the operating system and for the application program
if present is preferably securely stored in the solid state memory, e.g. in an encrypted
15 and/or scrambled form. The system can preferably not be booted from the disc
memory. Preferably, the server card has a serial bus for monitoring functions and states
of the server card.

Preferably, the server card is pluggable into a connector. Each server unit is
preferably pluggable into a local area network (LAN) on the server module which
20 connects each server to an administration card in the server module. A plurality of
server units are preferably connected via a connector into which they are pluggable to
a hub which is part of the server module LAN. A proxy server is preferably included as
part of the server module LAN for providing proxy server facilities to the server units.
Preferably, two proxy servers are used to provide redundancy. Access to the LAN of
25 the server module from an external network is preferably through a switch which is
included within the LAN. The server module may be located in a local area network
(LAN), e.g. connected to a switch or in a wide area network, e.g. connected via switch
with a router or similar.

The server module preferably has a local management system capable of
30 receiving a remote command (e.g. from a network) and executing this command so as
to modify the service performance of at least one of the server cards. Modifying the
service performance means that more than just reporting the state of a server card or
selecting a server card from the plurality thereof. That is the local management system

is capable of more than just monitoring the server cards.

The server module may also include a load balancing unit. This load balancing unit may be used for load balancing applications running on the servers, e.g. an application may be provided on more than one server unit and the load on each server unit within the group running the application is balanced by the load balancing unit; or a load balancing unit may be used for load balancing network traffic, e.g. to balance the loads on proxies used to transmit received messages to the relevant server unit.

The present invention also includes a digital processing engine mounted on a card, for instance to provide a server card, the card being adapted to be pluggable into a connector, the digital processing card comprising:

a central processor unit; and a first rewritable, non-volatile disk memory unit mounted on the card. The engine is preferably a single board device. The digital processing card may also include a second rewritable solid state memory device (SSD) mounted on the card. The SSD may be for storing an operating system program and at least one application program for execution by the central processing unit. The card may be adapted so that the central processor is booted from the solid state memory device and not from the rewritable, non-volatile disc memory unit. Preferably, the disk memory is a hard disc. Preferably, more than one hard disk is provided for redundancy. An input/output device may also be mounted on the card. For example the I/O device may be a communications port, e.g. a serial or parallel port for communication with the CPU. The card is preferably flat (planar) its dimensions much such that its thickness is much thinner than any of its lateral dimensions, e.g. at least four times thinner in its thickness than any of its lateral dimensions. The processing engine preferably has a bus connection for the receipt and transmission of management messages.

Whereas the current ASPR technology is based on *IT Centric* platforms, one aspect of the present invention is an ASPR *network centric* environment. The provisioned applications would be offered under a subscription format to potential subscribers that would be able to "consume" the applications rather than acquiring the applications prior to their usage. In essence, the application consumer (e.g. e-merchant, e-businesses or e-university) would be liberated from the financial and technical burden that comes with acquiring and installing new applications and keeping those applications up-to-date.

Application customers and users benefit from the economies of scale for the

shared infrastructure, but also expect high service levels and predictable costs for their business critical applications. As concrete examples, the data transmission is increased, the security level is higher, and the prices are conveniently packaged with fixed monthly payments. The present invention may be deployed by Application Service
5 Providers (ASPR). Application service provisioning is provided in which application software is remotely hosted by a third party such as an ISP (Service provider in general) that is accessed by the subscribing customer over the (Internet) network.

The present invention will be described with reference to the following drawings.

10

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic representation of a conventional wide area data carrier network.

15 Fig. 2 is a schematic representation of a conventional wide area data carrier network in accordance with an embodiment of the present invention.

Fig. 3 is a schematic representation of a server module in accordance with an embodiment of the present invention.

Fig. 4 is a schematic representation of a server chassis in accordance with an embodiment of the present invention.

20 Fig. 5 is a schematic representation of a management chassis in accordance with an embodiment of the present invention.

Fig. 6 is a schematic representation of a server card in accordance with an embodiment of the present invention.

25 Fig. 7 is a schematic representation showing how the proxy server of the management chassis transfers requests to an individual server card in accordance with an embodiment of the present invention.

Fig. 8 is a schematic representation of a how the configuration is uploaded to a server card on boot-up in accordance with an embodiment of the present invention. The management database contains configuration details of each server card.

30 Fig. 9 is schematic representation of how management information is collected from a server card and transmitted to a remote operations centre in accordance with an embodiment of the present invention.

Fig. 10 is a schematic representation of a server module in accordance with an

embodiment of the present invention used in a local area network.

DESCRIPTION OF THE ILLUSTRATIVE EMBODIMENTS

5 The present invention will be described with reference to certain embodiments and to certain drawings but the present invention is not limited thereto but only by the claims. For instance a wide area network will be described with reference to wireline telephone access but the present invention is not limited thereto and may include other forms of access such as a Local Area Network, e.g. an Intranet, a Wide Area Network, a Metropolitan Access Network, a mobile telephone network, a cable TV network.

10 One aspect of the present invention is to provide server capability in premises which can be owned and maintained by a the telecom provider, for example in a "point-of-presence" (POP) 12. Another aspect of the present invention is to provide a Remote Access IP network infrastructure that can be deployed anywhere in a wide area network, for example, also at the edges of the network rather than exclusively in a
15 centralised operations centre. Yet another aspect of the present invention is to provide a distributed server architecture within a wide area telecommunications network such as provided by public telephone companies. Yet a further aspect of the present invention is to provide a network management based architecture (using a suitable management protocol such as the Simple Network Management Protocol, SNMP or
20 similar) to remotely configure, manage and maintain the complete network from a centralised "Network Management Centre" 10. The SNMP protocol exchanges network information through messages known as protocol data units (or PDU's)). A description of SNMP management may be found in the book "SNMP-based ATM network Management" by Heng Pan, Artech House, 1998. From a high-level
25 perspective, the message (PDU) can be looked at as an object that contains variables that have both titles and values. There are five types of PDU's that SNMP employs to monitor a network: two deal with reading terminal data, two deal with setting terminal data, and one, the trap, is used for monitoring network events such as terminal start-ups or shut-downs. Therefore, if a user wants to see if a terminal is attached to the
30 network, SNMP is used to send out a read PDU to that terminal. If the terminal was attached to the network, the user would receive back the PDU, it's value being "yes, the terminal is attached". If the terminal is shut off, the user would receive a packet sent out by the terminal being shut off informing of the shutdown. In this instance a

trap PDU would be dispatched.

The deployment of the equipment in the network edges, for example in the POPs 12, can be done by technicians because the present invention allows a relatively simple hardware set up. The set-up is completed, e.g. configuration and security set up, by the network engineers remotely via the network, e.g. from a centralised operations centre 10. If modifications of the structure are needed, this can usually be carried out remotely without going on-site. If infrastructure changes or upgrades in the network edges are mandatory, such as increasing incoming line capacity, technicians can execute the required changes (in the above example by adding network cards) whilst the network engineers are monitoring remotely the progress and successful finalisation.

An embodiment of the present invention will be described with reference to Fig 2. A wide area network 1 which may span a town, a country, a continent or two or more continents is accessed by an access network 2. The access network will typically be a wireline telephone or a mobile telephone system but other access networks are included within the scope of the present invention as described above. Typically, POP's 12 are placed at the interface of the wide area network or data carrier network 1 and an access network 2. Application servers installed in server modules 20 may be located anywhere in the data carrier network 1, for instance, in a network POP 12, 14 whereas management of the server modules 20 is achieved by network connection rather than by using a local data input device such as a keyboard connected to each server. For instance, the server modules 20 may be located at the edges of the network 1 in the POP's 12 and are managed centrally through a hierarchical managed platform 16 in an operations centre 10, e.g. via a suitable management protocol such as SNMP. Preferably, application servers (e.g. for providing applications for e-Shops, E-Mail intranet servers, e-Game servers, Computer Based Training packages) are card mounted and are insertable as required in a standard chassis, for instance a plurality of chassis' can be installed in a standard 19" cabinet. The applications running on the servers 20 are preferably provisioned remotely. The application server module 20 runs server software to provide a certain service, e.g. a homepage of an e-merchant. The person who makes use of the server module 20 to offer services will be called a "user" in the following. A user may be a merchant who offers services via the Internet. The person who makes use of the server module 20 to obtain a service offered by a user, e.g. by a merchant, will be called a "customer".

A customer 11 can access the application running on one of the servers 20 located at the relevant POP 12, 14 from their own terminal, e.g. from a personal computer linked to an analog telephone line through a modem. In addition, each server of the group of servers in a server module 20 in a POP 12, 14 is remotely addressable, e.g. from a browser running on a remote computer which is in communication with the Internet. For instance, each server in a server module 20 has its own network address, e.g. a URL on the World-Wide-Web (WWW) hence each server can be accessed either locally or remotely. However, to improve graceful capacity upgrading and scalability, it is preferred if a server module 20 has a single Internet address and that each server in the server module 20 is accessed via a proxy server in the server module 20 using URL extensions. Thus, a server module 20, in accordance with one implementation of the present invention can provide an expandable "e-shopping mall", wherein each "e-shop" is provided by one or more servers. The server module 20 is remotely reconfigurable from an operations centre 10, for instance a new or updated server program can be downloaded to each of the servers in the server module. Each server of the server module 20 can also be provisioned remotely, e.g. by the user of the application running on the respective server using an Internet connection. This provisioning is done by a safe link to the relevant server.

Embodiments of the present invention are particularly advantageous for providing access to local businesses by local customers 11. It is assumed that many small businesses have a geographically restricted customer base. These customers 11 will welcome rapid access to an application server 20 which is available via a local telephone call and does not involve a long and slow routing path through network 1. The data traffic is mainly limited to flow to and from the POP 12 and does not have to travel a considerable distance in network 1 to reach a centralised data centre. More remote customers 11 can still access server module 20 and any one of the servers therein via network 1 as each server is remotely accessible via an identification reference or address within the network 1.

Even when a server module 20 is located in an operations centre 10 in accordance with the present invention, its provisioning and configuration is carried out via a network connection. That is, normally a server has a data entry device such as a keyboard and a visual display unit such as a monitor to allow the configuration and provisioning of the server with operating systems, server applications and application

in-line data. In accordance with the present invention all this work is carried out via a network connection, e.g. via a LAN connection such as an Ethernet™ interface.

The present invention may also be used to reduce congestion due to geographic or temporal overloading of the system. The operator of network 1 can monitor usage, for example each server of a server module may also provide statistical usage data to the network 1. Hence, the network operator can determine which applications on which server modules 20 receive a larger number of accesses from remote locations in comparison to the number from locations local to the relevant POP 12, i.e. the network operator can determine when a server application is poorly located geographically. This application can then be moved to, or copied to, a more suitable location from a traffic optimisation point of view. Applications can be duplicated so that the same service can be obtained from several POP's 12, 14. For instance, if a TV commercial is to be run which is likely to result in a sudden flood of enquiries to a particular server, the relevant application can be provisioned remotely on a number of servers located in server modules 20 in different geographic areas before the commercial is broadcast. The distributed access will reduce network loads after the broadcast. Thus, the present invention allows simple and economical scalability both from the point of view of the network operator as well as from that of the user or the customer.

A server module 20 in accordance with an embodiment of the present invention is shown schematically in front view in Figs. 3 and 4. A standard 19" cabinet 22 contains at least one and preferably a plurality of chassis' 24, e.g. 20 chassis' per cabinet 22. These chassis 24 may be arranged in a vertical stack as shown but the present invention is not limited thereot. Each chassis 24 includes at least one and preferably a plurality of pluggable or insertable server cards 26, e.g. 12 or 14 server cards in one chassis, resulting in a total of 240 to 280 server cards per cabinet 22. The server cards 26 are connected to an active back plane. In addition a management chassis 28 may be provided, e.g. one per cabinet 22, which is responsible for managing all the server cards 26 and for providing a remote management (for example, SNMP) and proxy server functionality. The management chassis 28 includes a switch 32 which is preferably extractable and a suitable interface 33 to provide access to the network to which the server module 20 is connected. The management chassis 28 may, for instance, be composed of 4 server cards 34-37, a patch panel 38 and a back plane 40 for concentrating the connection of the patch panel 38 and of the server cards 34-37.

The four server cards include at least one proxy server card 35, an optional proxy server card 37 as back-up, a load balancing card 36 and an administration card 34. The management chassis 28 is used to concentrate network traffic and monitor all equipment. The server cards 26 are interconnected via the patch panel 38 and one or more hubs 42 into a Local Area Network (LAN). This specific hardware solution meets the constraints of a conventional telecom room:

- Space: High density with 240 to 280 servers in a standard cabinet.
- Low heat dissipation, EMC Compliance.
- High availability through critical elements redundancy.
- Optimised Maintenance with easy access and removal of all components.

A chassis 24 is shown schematically in a top view in Fig. 4. It includes a plurality of server cards 26 plugged into a backplane 40 which is integrated with an active or passive hub 42. One or more power supplies 44 are provided for powering the server cards 26 and the hub 42 if it is in an active hub. The power supplies 44 are preferably hot swappable in case of failure. To provide cooling one or more fans 46 may be provided. Again, the fans 46 are preferably hot swappable. Each server card 26 is preferably planar with a connector for plugging into a back plane along one edge. The server card is preferably thin, e.g. its thickness should be at least four times less than any of its planar dimensions.

A management chassis 28 is shown schematically in top view in Fig. 5. It includes a plurality of printed circuit boards 34-37 plugged into a backplane 40 which provides a data bus as well as power connections. The printed circuit cards 34-37 may be of the same hardware design as the server cards 26 but are installed with different software. An extractable multi-media switch 32 is provided which is coupled to the server cards 26. Fans 46 and power supplies 44 are also provided.

Each server card 26 includes a server which has been stripped down to absolute essentials in order to save space and to lower power usage and heat generation. Each server card 26 is preferably pluggable so that it can be easily removed and replaced without requiring engineer intervention nor the removal of connections, wires or cables. A server card 26 in accordance with an embodiment of the present invention is shown schematically in Fig. 6. The components of server card 26 are preferably mechanically robust so that a card may be handled by technicians and not by specially

qualified engineers, e.g. without having using any other precautions than would be expected of a person inserting a memory card, a battery or a hard drive into a lap-top computer. The skilled person will appreciate from Fig. 6 that the server card 26 is configured to provide a programmable computer with non-volatile, re-writable storage.

5 Each server card 26 may include a central processing unit 52 such as an Intel Pentium™ Processor at 333Mhz, a random access memory unit (RAM) 54, e.g. 2 x 64 = 128 Mb of RAM memory, for example flash memory, a rewritable, non-volatile secure memory 55, e.g. a disk-on-chip memory unit 2000 M-Systems 11, a BIOS memory 53, e.g. a flash memory, and at least one rewritable, non-volatile storage
10 device 56 such as a hard drive or similar drive memory. Program code, e.g. the operating system as well as any system, network and server management programs are preferably included in the secure memory 55, e.g. encrypted and/or scrambled. User applications may be loaded onto the storage device 56 as would normally be done on a personal computer or a server, e.g. on the disc drive 56, however it is particularly
15 preferred in accordance with an embodiment of the present invention if each server card 26 is dedicated to a single user application. For instance, a specific application program or suite of programs is loaded into memory 55 to provide a single application functionality for the server card 26. This reduces the size of the memory 55 and simplifies operation of the server card 26. Preferably, this application program or suite
20 of programs is not stored on the hard drive 56 but is pre-installed into the memory 55. The hard drive 56 is preferably only used to store the in-line data necessary for a pre-installed program (e.g. in the case of an e-merchant program: colours of displays, prices, pictures of goods offered, video data, intilialisation parameters). Hence, each server card 26 preferably contains a solid state memory device (SSD) 55 which contains all
25 the software programs needed to run and control the application chosen for card 26. All the variable information such as user files and temporary files will be stored on the mirrored hard disks 56. Each hard disk 56 may be divided into at least two partitions, one being reserved for temporary files, log files and all system files which must be written. The system preferably contains two hard disk 56 which will be kept identical
30 through a mirroring/stripping mechanism, so that if one of the disks 56 fail the system stays fully operational. The two rewritable, non-volatile storage devices 56 may be two IDE hard disks of 10 Gbytes.

The isolation of system and user code from the storage device 56 (which can be

accessed by customers) improves security. Preferably, the storage device 56 is replaceable, i.e. pluggable or insertable without requiring complex or intricate removal of wiring or connectors. Such a replaceable storage unit is well known to the skilled person, e.g. the replaceable hard disc of some lap-top computers. Each storage device
5 56 is preferably mechanically held in place on the card 26 by means of a suitable clipping arrangement. The storage device 56 co-operates with the CPU 52, i.e. it is accessed after boot up of the processor unit 52 for the running of application programs loaded into memory 55. To allow communication with the LAN, at least one network interface chip 58 is provided. Preferably, two interface chips 58, 58' are provided, e.g.
10 two Fast-Ethernet™ 100 Mb interfaces. Also one serial bus connection (SM-bus 57) for the management of the server card is provided which is connected to the administration card 34 via the server module LAN. The SM-bus 57 carries management information, e.g. in accordance with the SNMP protocol.

A front panel 60 is provided with an RJ-45 jack for on-site monitoring purposes
15 via a serial communication port driven by a suitable input/output device 51 as well as an off-on control switch 64 and control indicators 66, e.g. LED's showing status, for instance "power off" or "power on". The server card 26 is plugged into a backplane 40. For this purpose the server card 26 includes a connector 68 which may be a zero insertion force (ZIF) connector. The backplane connection is for providing power both
20 to the server electronics as well as to the warning lights 66 on the front panel 60, as well as for connections to two fast-Ethernet 100 Mb connections 58, 58' and the one serial connection 57 for physical parameters monitoring. The fans 46 draw air from the back of the chassis 24. The air flow is designed to pass over the storage devices 56, which are at the back. The air passes over a heatsink on the CPU 52, which is located
25 towards the front.

The skilled person will appreciate that the server 26 provides a digital processing engine on a card which has all the items necessary to operate as such except for the power units. Thus an individual card may be plugged into a suitable housing with a power supply to provide a personal computer. Hence, the server card 26 may be
30 described as a digital processing engine comprising a disk memory unit 56 mounted on a motherboard.

The installation and operation of a server module 20 will now be described. A server module 20 comprises a number of server cards 26 installed into one or more

chassis' 24 and a management chassis 28 all of which are installed in a cabinet 22 and located in a POP 12. Each server card 26 is pre-installed with a specific application, although not all the server cards 26 must be running the same application.

The server module 20 includes a proxy server 35, 37 connected to the wide area
5 network 1 and is provided with remote management (from the operations centre 10) via a suitable management connection and protocol, e.g. SNMP version 1 or 2. The proxy server 35, 37 is preferably connected to the network 1 via a traffic load balancer. If the server module 20 is to be used with an Internet TCP/IP network, the proxy server 35, 37 may use the HTTP 1.1. protocol. Each server card 26 has a preinstalled
10 application which can be accessed, for example, by a customer browser. The configuration details of the home page of any server card 26 are downloaded remotely via the user who has purchased or rented the server card use. This information is downloaded via access network 2, e.g. by coupling a user personal computer or workstation to the respective server card 26 via the access network 2. Each user
15 prepares a command file using proprietary software which is transmitted to the relevant server card 26 in a safe messaging session protected by suitable authentication routines and encryption. All communications done between the user software and the server module 20 whatever the direction are encrypted using a suitable secure messaging system such as the Secure Socket Layer (SSL).

20 The proprietary software only needs to relate to the specific application for which the relevant server card 26 is dedicated and may include a software library especially designed for the user, e.g. for e-commerce. Once installed and provisioned, each server card 26 can be accessed remotely by either the user or a customer 11. Each server card 26 will be dedicated to one organisation (= one user) and to one
25 application and will not be shared between organisations. This increases security. In use, each server card 26 is monitored remotely via the network side management connections (SNMP) of server module 20. If a component defect is reported, e.g. loss of a CPU on a server card, a technician can be instructed to replace the defective card 26 with a new one. Such a replacement card 26 may have the relevant server
30 application pre-installed on it in advance to provide seamless access. If a hard drive 56 becomes defective, the stand-by hard drive 56 of the pair may be substituted by a technician. The load balancing card 36, the proxy server cards 35, 37 and the administration 34 may all have the same hardware design as server card 26. However,

the software loaded into memory 55 on each of these cards 34-37 is appropriate for the task each card is to perform.

A typical access of a server card 26 will now be described. On power up, each server card 26 boots using the content of the SSD 55 and will then configure itself, asking a configuration which it access and retrieves from the administration card 34. Since each server card 26 hosts a specific user it is mandatory that a card 26 is able to retrieve its own configuration each time it starts. The proxy server functionality is composed of at least two, preferably three elements. For instance, firstly the load balancing card 36 which distributes the request to one of the two proxy servers 35, 37 and is able to fall back on one of them in case of failure, e.g. if the chosen proxy server 35, 37 does not react within a time-out. Secondly, at least one HTTP 1.1 proxy server 35, 37, preferably two to provide redundancy and improved performance. Where redundancy is provided the load balancing card may be omitted or left redundant.

The procedure is shown schematically in Fig. 7. A customer 11 accesses the relevant WWW site for the server module 20. The network service provider DNS connects the domain with the IP address of the server module 20. The request arrives (1) at module 20 from the network at the switch 32 which directs (2) the request to the load balancing card 36 of the management chassis 28. The load balancing card 36 redirects (3, 4) the request to one of the two proxy servers 35, 37 depending upon the respective loading of each via the switch 32. The relevant proxy server 35, 37 analyzes the HTTP 1.1 headers in the request and redirects (5) the request to the right server card 26 using an internal IP address for the server card 26. This internal IP address of each server card 26 is not visible outside the server module 20. The server card 26 processes the request and sends (5) the answer back to the proxy server card 35, 37 which forwards the answer to the requester. This procedure relies on the HTTP 1.1 proxy solution. This means that the request will be redirected according to the domain name of the request. This information is provided by the HTTP 1.1 protocol. All 4.x and higher browsers (e.g. as supplied by Microsoft or Netscape) use this protocol version.

In order to avoid the typical updating problem with distributed processors, e.g. software maintenance and updating, on the server cards 26, a centralized network-side management of all parameters is implemented. If needed (upgrade of the server

application, security patch, load of virus signatures for an antivirus program) the administration card 34 is able to upload a new SSD (solid-state disc) image onto any or all of the server cards 26 and can force an upgrade of the system software. Any new boot scripts will also support all the automatic raid recovery operation upon the replacement of a defective hard disk 56. The administration card 34 is updated/managed as necessary via the network 1 from operations center 20. When a server card 26 boots, it retrieves its configuration from the administration card 34 (Fig. 8). First it retrieves its IP configuration according its position in the server module 20. Then it downloads all its configuration files and upgrades its software if needed.

10 Suitable protocols are used for these actions, e.g. DHCP (Dynamic Host Configuration Protocol) may be used for the IP configuration retrieval and TFTP for the software configuration. The DHCP solution will rely on the identification of the card by its MAC address (boot like). The updating procedure is therefore in two steps: firstly, an update is broadcast via network 1 to one or more server modules 20 where the update is stored in the administration card 34. Then on power-up of each server card 26, the update is loaded as part of the automatic retrieval procedure from the administration card 34.

When a server card 26 is assigned to a user it will be provided with its internal IP address. The server 20 allows basic monitoring and management through an HTML interface in order to allow decentralised management from the operations centre 10. This monitoring will be done through authenticated SSL connection (Secure Socket Layer protocol which includes encryption for security purposes). As part of the management function the server module 20 management data is transferred to the operations centre 10 in accordance with (MIB) Management Information Base

20 II. In addition it is preferred to extend this protocol to allow additional states to be monitored, e.g. a MIB II+ protocol, for recording and transmitting additional events as well as data useful to the provider of network 1 such as network utilisation. The MIB II Enterprise extension is provided to allow the monitoring of each server card 26 of a server module 20. Information about the configuration, the running status, network statistics may be retrieved. Physical parameters such as fan speed,

30 temperature, of each chassis 24 may also be monitored remotely by this means. The monitoring may be performed by a sequence of agents running on the relevant part of the system, e.g. an SNMP agent 72 responsible will collect or set information from

configuration files, will get real time statistics from each server card 26 and will get data from physical sensors in the chassis' 24. Preferably, a middle agent 74 monitors all SNMP traps, pool statistics from the server cards 26 and will be able to react to specific errors and transmits these to the remote operations centre 10 via network 1
5 (Fig. 9).

The management system provided in the management chassis 28 (MCH) allows a telecommunications network operator to manage a full cabinet 22 and up to 20 chassis 24 as a standalone equipment and to deliver high-added value services with QoS definition, to customers and users. From the point of view of the network
10 management centre 10 a server module 20 is seen as one network equipment with its own network addresses and environment.

"Service" may be understood as a synchronous group of applications running on "n" servers 26 (with assumption that n is not null). The application can be User (or Customer) oriented or System oriented. (User oriented applications can be a web
15 hosting or e-commerce application, for example. A System oriented application can be a Proxy, a Core Administration Engine, for example)

Customers of services delivered by in accordance with the present invention can access these services using a dedicated Service ID (SID). Moreover even if a service is accessible through a unique SID, it can be hosted on several servers. This is
20 possible using proxy solutions.

A "proxy", may be seen, for example as a piece of software, for example an object allowing entities to communicate together through a unique point. The proxy is able either to split, to copy, or to concentrate the network communication according specific rules. These rules are typically based on Layer 2 to Layer 7 protocol
25 information. The proxy can also change the nature of information it uses by translating it in order to match the needs of the involved entities, for example protocol conversion. A proxy may collect information, e.g. management information, or receive this information from one or more of the servers. A proxy may therefore allow monitoring and control, protocol conversion, may implement access control and may also co-
30 ordinate or manage several objects, e.g. applications running on several servers.

Four main management sub-systems can be identified:

An administration sub-system which allows remote administration and monitoring.

A processing sub-system allows an appliance such as a server to provide a service

A storage sub-system which allows a server to store data.

A synchronization sub-system which is dedicated to a storage sub-system and a
5 processing sub-system. It allows data replication over several servers and make applications synchronous.

If a service is hosted on several server cards 26, each server card 26 can process the whole request by itself. Nevertheless, if some modification of data is needed, all servers in a group must be synchronized. If an action is leading to data modification,
10 the server responsible for this operation will update all other servers in its service to maintain the data synchronized. Each server will access to its data through a “data proxy”, which will locally resolve the consultation of data, and will replicate all changes over all the servers hosting the service.

A service can be a Management Service (MSV) or a User Service (USV)
15 depending on the nature of the application. This service is accessible by its SID (typically a domain name or some protocol specific information: socket number, protocol id, etc).

Management Services (MSV) are hosted in the management chassis MCH 28 and they provide functionality to all the other services. For example “Administration
20 service” or “SSL Proxy service” are MSV.

MSV typically can be classified in two families:

- MSC, the Management Services Communication oriented which include all MSV that directly allow communication between customers or users and user services USV. An example is a Proxy Service or a Load Balancing Service which allow the
25 making of the link between the customer or the user and the service through the network name of the server module 20.
- MSH, the Management Services Help oriented which include all MSV that provide intermediate services, or help, to other MSC. For example a service which can provide, store, or monitor information about the others services is a MSH.

30 A User Service (USV) provides a service to a customer or a user. Typical USV are Web Hosting Application, e-Shop Application or e-Mail Application.

USV can be implemented in two major configurations:

- When the focus is reliability, the service is delivered by an application running on 2 servers, one backing up the other.
- When the focus is on performance, the service is delivered by an application running on n servers (in this case $n > 1$) combined with an application load balancing service allowing repartition of the load between the servers. A side effect of this solution is an improvement of the reliability.

The Load Balancing Service (LBS) is used to balance requests on several server cards 26 according to a specific algorithm. These server cards host the same application and the LBS allows these servers and their applications to deliver a specific service. The LBS can be hosted on up to two servers allowing high availability. To reach a service, a customer or user will need an SID that can match the name of the service to the network address of the corresponding server module 20. For example with IP-based applications, this external name server is a domain name server (DNS); other type of directories can be used, however. With this network address, the user or customer will be able to reach the server module 20 through the network. This network address is bound to the "access proxy service". The proxy will find the internal network address of the service, extracting information from protocol-determined fields in order to achieve the internal routing of the request. This routing done, a communication channel is opened between the user or customer and the service. All proxy services are designed to work on behalf of an LBS.

If the service is running on several servers, several implementations are included within the scope of the present invention:

- the proxy service can select the server card 26, which will process the request, according to several parameters, e.g. load on the server cards, availability, cost of access.
- the proxy service can select the network address for the service. One server card in the service group owns this address, if this server card fails another in the service group will take the ownership of the address.

It is possible to proxy all protocols if it is possible to extract from the protocol any information allowing to direct communication to the right service. Different types of proxy service which may be used with the present invention are:

HTTP Proxy: The HTTP Proxy service allows binding a URL identification

with an internal IP address used as locator in the server module 20.

SSL proxy: The SSL Proxy service allows to provide SSL based ciphering for a whole server module 20. A dedicated DNS name is given to the server module 20.

Though this specific naming an application can accept secured connection.

- 5 FTP Proxy: The FTP Proxy service allows exchanging files according to the FTP protocol. A user will be able to send or receive file to/from its service through the server module network address and a personal login.

- 10 POP3 Proxy: The POP3 proxy service allows to access mailboxes according to the POP3 protocol. A user will be able to receive e-mails from its service through the server module network address and a personal login.

The Administration Service (ADS) allows management of a full cabinet like a single network equipment.

This management can be performed using three different interfaces:

- SNMP V1 and V2 with trap notifications
- 15 • HTML Interface
- Shell Interface (Command Line Interface)

- All management interfaces are connected to the Core Administration Engine (CADE) through a specific API. The CADE maintains the configuration and the status of all components of a server module 20. These components are running software, hardware
- 20 and environmental parameters. Each server card 26 can communicate with the CADE as a client/server and the CADE can communicate with all servers in the same way.

Each server is running software dedicated to answer CADE. This software can:

- (i) Introduce an application or server card 26 to the ADS and check its status (provisioned or not for example).
- 25 (ii) Notify ADS asynchronously with failures or exceptions
- (iii) Answer to hardware/software status request
- (iv) Run predefined actions for remote control.

- The communication protocol used between the CADE and the server cards 26 is not depending on the nature of managed application. The ADS can be hosted on
- 30 two server cards. One in backup of the other to improve reliability and availability of this service.

The ADS maintains information about each server card 26 and each service

over the complete cabinet it manages. Information about services is relevant to running the service and to service definition.

For each server card 26 the ADS stores, for example:

- All log files
- 5 • Descriptor of hardware installed on the server card (stored at each server boot)
- Descriptor of software installed on the server card and their releases/versions (checked/updated at each server boot)
- Software running on the server card (stored at each request or on asynchronous notification)
- 10 • Status of the server card (Running, Off, Assigned, Available)
- Environmental parameters (network interface status, disk status, memory, disk and CPU load, etc.)

For service definitions:

- Name and definition of software which is installed on a server card to perform
15 this service
- Description of possible actions on the software and available status
- Plug-ins to generate configuration file from the service definition

For the active services:

- The public name of the service
- 20 • A list of parameters defining the service levels
- The server cards involved by the service and their role (Service Load Balancing, Master, Backup, etc.)
- Statistics of accesses (number of hits per timeframe, load, etc.)
- Security parameters such as Access Control Lists (ACL)

25 Monitoring in a server module can be performed in two ways:

- Asynchronous notification: each server card is monitoring its services and sends an alert if something goes wrong.
- Polling: ADS monitors hardware status of a chassis by polling each elected server card 26 in the chassis and the ADS checks status of running server cards 26.
- 30 • Monitoring is also used to feed information into a database.

Other available services are, preferably:

- Billing Service (BS): collects all information about bandwidth, time and

resource usage needed for accounting activities.

- Performance Reporting Service (PRS): allows users of the services to obtain measurement of the QoS they have subscribed for.
 - Secured Payment Gateway Service (SPGS): provide to a server module
- 5 with a payment gateway for all on-line payment needs linked to e-commerce appliances.

A USV provides a service to a user and/or a customer. Moreover, each USV may be associated with a dedicated Application Load Balancing Service (ALBS). This service, which is similar by nature to the MLB service, allows a load balancing of all

10 requests to the USV between the server cards hosting this service. A USV is not linked to a specific software, it is a set of software allowing the provisionning of a high value customer oriented service.

Provisioning an USV consists in binding a server card 26 or a group of server cards 26, with an ID, service levels and credentials. As soon as the provisionning

15 phase is completed, the server module 20 is ready to deliver the service.

The main phases in the provisionning procedure are:

1. Provisioning parameters through one of the administration interfaces.
2. The CADE binds the service to a server card or server cards. The number of server cards involved and their location are determined by the service level parameters.
- 20 3. The ADS prepares the configuration of the relevant software to be started using specific plug-ins.
4. Once server cards 26 are chosen and configuration files are ready, the CADE communicates with all involved server cards in order to setup each server card to provide the service according to given parameters.
- 25 5. Then ADS notifies the proxy service that the new service is available.

When the Network Management Centre (NMC) 10 decides to apply an update to the different server modules 20, the mechanism is based on a two-step update mechanism. The update will generally contain software updates or patches. An update is contained in one file, with a specific format, that does not only contain information,

30 which must be upgraded, data are packed with specific upgrade software that is able to apply updates according to versioning information installed on each server card 26.

The versioning system and the build process automatically generate this software. The

upgrade software generated allows migration between one build to another. This upgrade software is responsible of backing up every data it may change and generating the "downgrade scripts" in order to reverse the upgrade in case on failure. It may also include a data migration module in order to upgrade the storage schemes.

- 5 All information needed by the ADS to manage its matrix of server cards 26 is stored in a database that mainly contains, per server card, configuration files, SLA profiles, application profiles and system descriptor. An example of the entries in the database are given below.

	Server card I-Brick 2.1	Server card I-Brick 2.2
10	<ul style="list-style-type: none"> •Name www.sell.com •IP Address10.1.2.1 •Config Path/iboot/2.1 	<ul style="list-style-type: none"> •Name www.sell.com •IP Address10.1.2.2 •Config Path/iboot/2.2 <ul style="list-style-type: none"> •Hosts •Passwd •Shadow •Apacheconf •Snmpconf •Squidconf •Network.conf •Quotasconf
15		

The update mechanism is as follows:

1. NMC 10 makes available the update/patch. This update/patch can be stored in
20 an optional common repository.
2. NMC 10 notifies different server modules 20 via the wide area network that this update/patch is available and must be applied to a specific profile within a specific time scale.
3. Each ADS uses its management database to select all involved server cards of
25 the server module 20 depending on the scope and the severity constraints attached to the update.
4. Each ADS manages the update distribution over its own server cards. That means the ADS controls and manages update/patch deployment on profile screening and can update either applications or operating system components including kernel on
30 each managed server card 26.
5. This mechanism is also available for the ADS itself in recurrent mode. In order to warranty availability of ADS, a protection mechanism may be implemented in order to monitor ADS processes and to restore the latest stable state for ADS in case of

trouble in the update process.

6. When the upgrade process is completed, the ADS notifies the NMC 10 with the new status.

The mechanism described above allows the NMC 10 to delegate all application and system update/patch operations to the different ADS embedded in all the server modules 20 deployed on the network. The associated side effect is an optimization of the bandwidth usage for this type of operations.

In the above, a server module 20 in accordance with the present invention has been described for use in a wide area data carrier network. The server module 20 as described may also find advantageous use in a Local Area Network as shown schematically in Fig. 10. For example, LAN 80 may be an Intranet of a business enterprise. Server module 20 is connected in a LAN 80. Server module 20 may have an optional connection 81 to a remote maintenance centre 82 via LAN 80, a switch 83 and a router 88 or similar connection to a wide area network, e.g. the Internet to which centre 82 is also in communication. The LAN 80 may have the usual LAN network elements such as a Personal Computer 84, a printer 85, a fax machine 86, a scanner 87 all of which are connected with each other via the LAN 80 and the switch 83. Each server card 26 in the server module 20 is preferably preinstalled with a specific application program, such as a text processing application such as Microsoft's WORD or Corel's WordPerfect, or a graphical program such as Corel Draw, etc. Each PC 83 can retrieve these programs as required – for each different application a different server card 26. In addition, a server card 26 may be allocated to each PC 84 for file back-up purposes on the hard disk 56 thereof. 240 to 280 server cards provide ample server capacity to provide a Small or Medium sized Enterprise with the required application programs and back-up disc (56) space.

In case one of the server cards 26 goes down, it is only necessary for a similar card with the same application to be installed. While other applications can continue running. This improves outage times of the system and increases efficiency. The loss of a server card 26 may be detected locally by observing the status lights on the front panels 60 of the server cards 26. Alternatively, the operation of server cards 26 may be monitored by the maintenance centre 82 as described above for operations centre 10. Also software updates may be sent from maintenance centre 82 in the two step updating procedure described above.

CLAIMS

1. A wide area data carrier network comprising:
one or more access networks;
5 a plurality of server units housed in a server module and installed in said wide area data carrier network so that each server module is accessible from the one or more access networks, each server module including a management system local to the server module for managing the operation of each server unit in the module; and
an operations centre for remote management of the server module via the local
10 management system, the server module being connected to the operations centre for the exchange of management messages through a network connection.
2. The wide area network according to claim 1, wherein the local management system is adapted to receive a management message from the operations centre containing a
15 command and for executing this command to modify the service performance of at least one server unit.
3. The wide area network according to claim 1 or 2, wherein the server units are active servers.
20
4. The wide area network according to any of claims 1 to 3, wherein the management messages comprise at least any one of: remote monitoring of the status of any server unit in a module, trapping alarms, providing software updates, activating an unassigned server module, assigning a server module to a specific user, extracting
25 usage data from a server module or server unit, intrusion and/or hacker detection.
5. The wide area network according to any previous claim, wherein each server unit includes a central processor unit and a secure memory device for storing the operating system and at least one application program for running the server unit.
30
6. The wide area network according to claim 5, wherein the secure memory device is a solid state device.

7. The wide area network according to any previous claim wherein each server unit comprises a rewritable, non-volatile disc storage device.

8. The wide are network according to claim 7, wherein the server unit is adapted so
5 that the rewritable, non-volatile storage device contains only data required to execute the application program and/or operating system program stored in the secure memory device but does not contain program code.

9. The wide area network according to claim 8, wherein the central processing unit not
10 bootable via the rewritable, non-volatile storage device.

10. The wide area network according to any previous claim wherein each server unit is mounted on a pluggable card.

11. The wide area network in accordance with any previous claim, wherein the server
15 module is located in a point of presence (POP).

12. A method of operating a wide area data carrier network having one or more access networks comprising the steps of:

20 providing a plurality of server units housed in a server module in said wide area data carrier network so that each server unit is accessible from the one or more access networks;
managing each server unit in a server module by means of a management system local to the server module;
25 additionally managing each server unit of the server module remotely through a network connection to the server module via the local management system.

13. The method according to claim 12, further comprising the steps of:
the local management system receiving a command through the network connection
30 from the data carrier network and executing the command to change the service performance of at least one of the server unit.

14. The method according to claim 12 or 13, wherein each server unit is pluggable, further comprising the step of removing a server unit from a server module and plugging a server unit into the server module.
- 5 15. A server module comprising:
at least one server card insertable in the server module, the server card having a central processing unit and at least one rewritable, non-volatile disk memory device mounted on the card.
- 10 16. The server module according to claim 15, wherein an Input/Output (I/O) device is mounted on the server card.
17. The server module according to claim 15 or 16, wherein at least one local area network interface is mounted on the server card.
- 15 18. The server module according to any of claims 15 to 17, further comprising a solid state memory device mounted on the server card.
19. The server module according to claim 18, wherein the operating system for the
20 central processing unit and optionally at least one application program is pre-installed in the solid state memory device.
20. The server module according to any of the claims 15 to 19, further comprising a proxy server.
- 25 21. The server module according to any of claims 15 to 20, further comprising a management unit for managing the server card.
22. The server module according to claim 21, wherein the local management unit is
30 adapted to receive a management message containing a command from external and for executing this command to modify the service performance the server card.

23. The server module according to any of claims 15 to 22, wherein the server card has a management bus connection.
24. A digital processing engine mounted on a card, the card being adapted to be pluggable into a connector, the digital processing card comprising:
a central processor unit; and
a rewritable, non-volatile disk memory unit mounted on the card.
25. The engine according to claim 24, further comprising a rewritable non-volatile solid state memory device (SSD) mounted on the card.
26. The engine according to claim 25, wherein the SSD stores an operating system program and at least one application program for execution by the central processing unit.
27. The engine according to any of claims 24 to 26, wherein the disc memory is a hard disc.
28. The engine according to any of claims 24 to 27, further comprising an input/output device on the card.
29. The engine according to any of claims 24 to 28, further comprising a management bus connection.
30. The engine according to any of claims 24 to 29, wherein the engine is a server.




Fig. 1




Fig. 2




Fig. 3






Fig. 4









Fig. 7



8
Lij




Fig. 9

